ISCR

Institute for Scientific Computing Research

Student Internship Research Summaries

CASC Center for Applied Scientific Computing

Multiresolution Geometry Visibility Determination for ROAM

Lucas Ackerman
Worcester Polytechnic Institute

Summary:

The Realtime Optimally Adapting Mesh algorithm (ROAM) enables Dynamic view-dependent level of detail rendering of multiresolution geometry. In order to reduce rendering overdraw and to allocate triangles only to visible surfaces, it is necessary to have a visibility determination algorithm that operates on multiresolution geometry and scales well to very large data sets.

The algorithm developed for this purpose utilizes an on-demand, dynamic occluder selection process, and can refine occluders as best fits the current viewpoint. The method of scene traversal is based on Frustum Advancement, testing only visible scene portions, and so has a small memory footprint even when dealing with very large geometry sets. To take advantage of frame-to-frame visible set coherence, the algorithm may use eye-cell or occluder-shrinking methods to determine visibility for a region of space around the eye point. The algorithm requires only a coarse octree constructed on base-mesh triangle volumes as a pre-compute step, and so can leverage fast visibility computation of complex scenes without costly computation or storage of potentially-visible-set structures. Finally, the algorithm can also generate a simple quadtree-based Z-buffer, which may be used to check visibility of mobile objects in the scene.

Loop Fusion Optimization for Array Statements with ROSE

Marcel Arndt
University of Bonn, Germany

Summary:

Computing power increased dramatically within the last decade, but performance optimization of numerical applications did not loose any of its importance. Many methods, including expression templates within object-oriented libraries, do a fine job, but are not completely satisfactory, since they are not able to exploit semantic coherences. This is where the C++ preprocessor ROSE, which is being developed within the Overture group at CASC, starts, by applying certain transformations to the source code before compilation. One sort of transformation is based on loop fusion of array statements with respect to cache structures. I developed methods to explore the optimal shape of these optimizations, which are the foundation for a further implementation within ROSE.

The transformed array statements are modeled by C programs, which are generated automatically by several pen- and shellscripts. This allows a hatch-like quasi-exhaustive evaluation of performance subject to more than ten different parameters including degree of dependence between the arrays involved, the amount of loop fusion, the number and sizes of the arrays, and so forth.

For a better understanding and analysis of the results, the data is formatted into clear tables and PostScript files with graphical presentation are generated. For the latter purpose, a C++ program has been developed using Overture.

The results show that the fusion of several loops into one single loop can increase the performance up to 20% because of caching effects, but going too far with loop fusion leads to a code which runs as much as five times slower. We also found that a careful reordering of the statements based on a dependence analysis allows fewer loops and faster codes.

It is planned to implement these transformations with ROSE. My results will serve as a guide for the precise definition of those transformations.

A Least-Squares Finite Element Method for 3-D Neutron Transport Problems

Travis M. Austin
University of Colorado, Boulder
Department of Applied Mathematics

Summary:

The application of a least-squares finite element technique to the 3D neutron transport equation is an interesting alternative to current control volume methods. The method relies on a properly scaled least-squares functional to develop a minimization principle. The scaling is designed to capture the correct behavior of the solution across parameter regimes. The minimization principle provides a solid mathematical framework for building all of the components of a multigrid solution procedure.

Using a finite expansion of spherical harmonics to represent the angular variable leads to a selfadjoint, positive definite system of moment equations. A conjugate gradient method with a block Jacobi preconditioner, which uses multigrid to invert the diagonal elements, is used to approximately solve this system of equations. For nearly all parameter regimes and all diagonal elements, standard multigrid is sufficient as a preconditioner. However, for certain parameter regimes, multigrid performs poorly on the diagonal elements associated with the first-order moments. A new multigrid algorithm is being developed that simultaneously addresses the three first-order moments.

As always, the goal of multigrid is to develop a scheme that has convergence factors independent of discretization size. When the "grad-div" term is dominant, the error reduction factor for the standard multigrid algorithm approaches unity quickly as discretization size approaches zero, which is unacceptable. The problem lies in the "hiding" of some of the oscillatory error terms at the low end of the spectrum. Smoothing works only on eigenvectors at the high end of the spectrum. Usually, these high frequency eigenvectors contain the oscillatory errors. For the given neutron transport operator, other oscillatory errors at the low end of the spectrum are left unattended. When the smoothing does not truly smooth, the coarse-grid correction becomes worthless as a component in the multigrid algorithm.

During five weeks spent at LLNL I worked on development of a 3D code that will tackle this problem of divergence-free error. Britton Chang and Barry Lee of CASC have already shown that the FOSLS approach works very well in two of the three important parameter regimes studied in neutron transport. From our existing 2D experience, we are confident of developing a 3D code that is robust with respect to all regimes encountered in practice.

Djehuty, LLNL's Next Generation Stellar Evolution Code

Brian Ball

Worcester Polytechnic Institute of Technology

Summary:

The Djehuty project is a next generation stellar evolution code. It will be capable of three-dimensional modeling of binary star systems. Since half of all 'stars' exist in binary systems, the Djehuty project will have strategic potential for the stellar simulation community. The Djehuty team at LLNL is comprised of people from many locations, including IGPP, CASC, A and V Divisions, and UC Berkeley. The code is derived from projects at LLNL that have been in existence for over three decades.

One of the main components is ARES, which is a multi-block, regularly connected, 2D and 3D multi-material radiation hydrodynamics code. One of my tasks was to help remove certain functions from the hydrodynamics code which are not applicable for stellar evolution. Another task was the addition and testing of a new equation of state, created by Peter Eggleton (IGPP), in the Djehuty code. Testing was done using a standard shocktube test problem with both the Eggleton equation of state and an ideal gas equation of state at the appropriate environmental conditions. Radiation diffusion was also implemented and tested in the Djehuty code by using a mock idealized spherical NIF capsule. The testing was done in serial and in parallel on the "Compass" Alpha cluster.

Adaptive Methods in Numerical Cosmology

Zackery Belanger
Oakland University

Summary:

The goal of this project is to investigate the use of adaptive mesh refinement (AMR) in the solution of systems of Einstein equations arising in cosmology. Numerical solutions of these systems are increasingly important aids to understanding the nature of singularities characterizing various theoretical models of the origin and evolution of the universe. The solutions of cosmological systems such as the Gowdy T3 model are known to develop localized fine-scale structures as they evolve in time. The efficient solution of such equations requires the ability to employ computational grids with high resolution only where it is needed.

Under the supervision of Xabier Garaizar and Milo Dorr in the Center for Applied Scientific Computing (CASC), I have developed a code to solve the Gowdy T3 system of equations using a Godunov algorithm. The Gowdy system has the form of a variable-speed wave equation with nonlinear source terms. The solutions of this system represent wave amplitudes in a space-time metric. An important issue that I had to address was the treatment of the nonlinear source terms. The development of an accurate and robust uniform grid integrator is the key step in creating an adaptive algorithm using block-structured, locally Cartesian grids. I have begun to incorporate my integrator in an adaptive code using the SAMRAI (Structured Adaptive Mesh Refinement Application Infrastructure) system under development in CASC.

I will continue to develop my adaptive code as part of my Master's thesis at Oakland University. The next goal will be to implement a higher-order version of the current Godunov integrator. I will then use the new code to investigate the adaptive solution of extended models, such as the magnetic and twisted Gowdy systems, which have not yet been numerically simulated, pending the availability of the resolving power of AMR.

Parallel Volume Rendering of Unstructured Data

Janine Bennett
University of California, Davis

Summary:

There is a need for ways to visualize large-scale data sets. Current methods can be very slow and inefficient. By parallelizing portions of HIAC (High Accuracy Volume Renderer), a legacy volume renderer for unstructured data, we can make the visualization of these large-scale data sets much quicker.

The portion of the HIAC code responsible for visibility sorting, polyhedron projection, and the OpenGL calls used to render to the screen is now parallelized using Pthreads and two buffers for inter-thread communication. One thread is used for visibility sorting, one is used for the OpenGL calls, and all remaining threads are used to project the polyhedron. We tested our Pthreads code on a 48-Processor SGI Onyx2 with 8 Infinite Reality pipes running IRIX 6.5.7. The parallelized version of code runs twice as fast as the serial version of the code. Peter Williams, Randy Frank and Deborah Walker were all involved in different stages of the project.

We are investigating ways to further parallelize by processing the input data into several load-balanced screen tiles. For each tile we will apply the threaded code to generate images even faster.

Symmetric Wavelets on Arbitrary Topology

Martin Bertram
University of California, Davis

Summary:

Wavelets are used in lossy and lossless compression schemes for scientific data obtained from flow field simulations on supercomputers. Compression is necessary to store, transmit, and visualize massive volume datasets. Most compression schemes are restricted to data defined on regular, rectilinear grids. For compact representation of isosurfaces and material boundaries of arbitrary topology, however, wavelet compression schemes need to be generalized to handle data defined on polygonal meshes.

We constructed new wavelet basis functions defined on subdivision surfaces that have a polygonal base mesh and generate subgrids of regular topology, such as Catmull-Clark subdivision surfaces. Compared to most wavelet approaches on arbitrary topology, our wavelets generalize tensor product basis functions. As a consequence, we obtain B-spline representations for reconstructed data, except in the neighborhood of a few extraordinary points (with valence not equal to four) in the base mesh.

We implemented two new wavelet transforms that generalize bilinear and bicubic B-splines to mesh domains of arbitrary topology. To obtain compression, we encode the wavelet coefficients that are sparse or have small absolute values using a coding scheme that we developed earlier. The overall compression algorithm is highly efficient and provides surface reconstructions at multiple levels of accuracy.

We will use the new surface compression scheme for efficient extraction and visualization of isosurfaces. We also want to develop similar techniques for compression of volumetric data defined on irregular polyhedral meshes.

Web Voting Application to Assist in Establishing Software Development Standards

Melvina Blackgoat
Northern Arizona University

Summary:

The development of scientific software by and for many parties increasingly requires a mechanism for reaching agreement between the collaborative development groups. There exists no known method of voting on software standards between the Common Component Architecture (CCA) and the Equations Solver Interface (ESI) groups that span several DOE laboratories. Voting and discussion is done during infrequent meetings or through an email based voting system which requires significant amounts of staff time. For collaborative software development to succeed, a more effective voting mechanism is needed to establish software development standards.

We are developing a distributed voting application to facilitate DOE wide convergence on component standards for scientific computing. The web voting application uses a three-tiered architecture: a WWW browser interface implemented using HTML, a middle tier using Java Servlet Pages (JSPs) and additional Java, and a MySQL server as the bottom tier. We began development by creating JSPs that enter forum data into a database. JSPs are a combination of HTML content and behavior coded in Java. The web voting application will allow users to easily access information, make comments and vote on a particular forum in a small amount of time.

This software will be used by the DOE Common Component Architecture (CCA) and Equation Solver Interface (ESI) standards forums to manage voting on software interface standards. Automated web-based voting will simplify multi-lab collaboration and development of new software standards.

Reconstruction of Material Boundaries from Data Sets with Volume Fraction Information

Kathleen S. Bonnell
University of California, Davis

Summary:

There are numerous instances in which it is necessary to reconstruct or track the boundary surfaces (or "interfaces") between multiple materials that arise in simulations. Multi-fluid Eulerian hydrodynamics calculations require geometric approximations of fluid interfaces to form the equations of motion to advance these interfaces correctly over time. This project presents a new algorithm for material boundary interface reconstruction from data sets containing volume fractions.

To solve this problem, we transform the reconstruction problem to a problem that analyzes the dual data set, where each vertex in the dual mesh has an associated barycentric coordinate tuple that represents the fraction of each material present. After constructing a dual tetrahedral mesh from the original mesh, we construct material boundaries by mapping a tetrahedron into barycentric space and calculating the intersections with Voronoi cells in barycentric space. These intersections are mapped back to the original physical space and triangulated to form the boundary surface approximation. This algorithm can be applied to any grid structure and can treat any number of materials per element/vertex.

In typical simulations, the grid cells contain fractional volumetric information for each of the materials. Each cell C of a grid S has an associated tuple ($a1, a2, ..., am$) that represents the portions of each of m materials in the cell, i.e., ai represents the fractional part of material i. We assume that $a1 + a2 + ... + am = 1$. The problem is to find a (crack-free) piecewise two-manifold separating surface approximating the boundary surfaces between the various materials. We consider the dual dataset constructed from the given data set, with each point having the associated tuple. Thus, the boundary surface reconstruction problem reduces to constructing the material interfaces for a grid where each vertex has an associated barycentric coordinate representing the fractional parts of each material at the vertex. We use this "barycentric coordinate field" to approximate the material boundary surfaces.

If we have a data set containing m materials, we process each tetrahedral cell of the grid and map our tetrahedral elements into an m-simplex representing m-dimensional barycentric space. Next, we calculate intersections with the edges of Voronoi cells in the m-simplex. These Voronoi cells represent regions where one material "dominates" the other materials locally. We map these intersections back to the original space and triangulate the resulting points to obtain the boundary.

Concerning future work, we would like to insert a "measure and adjust" feature into the algorithm. Once an initial boundary surface approximation is calculated, we calculate (new) volume fractions for cells directly from this boundary surface. This will enable us to calculate the difference between the original volume fractions and the volume fractions as implied by our initial boundary surface approximation. It is then possible to adjust our material interfaces to minimize the volume fraction deviations. We also plan to extend this algorithm to multidimensional grids.

Line-Implicit Time-Stepping for the Solution of Navier–Stokes Equations on Overset Grids

Lars Carlsson

Chalmers University of Technology, Sweden

Summary:

Boundary layers arise in many fluid dynamics applications. Under popular fully implicit global time-stepping schemes, the solutions are often not resolved in time, and interesting flow phenomena disappear. Instead, a fully explicit method can be used. However, for these methods the time step becomes very small because of the large velocity gradients in the boundary layer. Our goal has been to alleviate the severe time-step restriction without stepping over the interesting temporal behavior. This has been accomplished by using line-implicit time-stepping.

We have implemented a two-dimensional line-implicit time-stepping scheme within the Overture framework. Overset grids have been used to divide the computational domain, and different time-stepping schemes have been used in the advancement of the solution on different component grids. The line-implicit time-stepping is designed to be used in highly stretched near-wall grids, where boundary layers occur.

The Laplacian and other spatial operators in the momentum equations are divided into fast and slow time scales. In the time-stepping, the fast and slow components are treated implicitly and explicitly, respectively. The resulting nonlinear systems of discrete equations are banded. Each system contains the solution along one grid line. The computational cost is of the same order as for a corresponding explicit scheme. The stretching in the direction normal to the implicit direction does not influence the stability of the time integration. The time step is determined only by the eigenvalues associated with the tangential, explicit direction. Hence, the overall time step is greater than in the case of a fully explicit method.

Calculations on applications like two-dimensional wing sections demonstrate how the time-stepping scheme performs in practice.

Algebraic Multigrid based on Element Interpolation

Tim Chartier
University of Colorado, Boulder

Summary:

My research at LLNL is to apply Algebraic Multigrid (AMG) to problems of material elasticity with complex geometry. This work applies to energy research and nuclear weapons stockpile stewardship. To protect our environment, we must understand how these materials behave in storage over long periods. My research plays a role in efforts to create scalable solution algorithms for the PDEs that arise in such applications. In particular, I am researching coloring algorithms for Algebraic Multigrid solvers based on element interpolation (AMGe).

This summer I continued research on selecting a hierarchy of coarse grids that lead to effective intergrid transfer operators by exploiting knowledge of finite element connectivity. AMGe is a practical compromise between purely algebraic multigrid, which makes use only of the final assembled matrix, and methods that require a full geometric description of a PDE problem, such as agglomeration. Element information is frequently available in applications. I continued discussion with LLNL researchers whom are also investigating and testing ideas in this area. I worked closely with Jim Jones on both 2-level and multilevel coloring ideas in AMGe. Our current research focuses on developing a multilevel algorithm that maintains low element complexity, grid complexity, and convergence factors. It is difficult to balance these competing factors. My work has also involved discussions with Panayot Vessilevski, Van Henson, and Rob Falgout. Current research shows promise. Given current advancements, we hope to provide a fully algebraic multilevel AMGe algorithm in 2001.

Parallelization of the HIAC Volume Renderer

Richard Cook
University of California, Davis

Summary:

The HIAC volume renderer is a piece of software that performs volume rendering using unstructured grids. The cells are first sorted using variations of the MPVQ sort by Peter Williams. They are then projected them onto the viewplane using an algorithm by Nelson Max and displayed using OpenGL hardware rendering.

The projection step is the bottleneck in the algorithm. Therefore, in this project, the projection and rendering steps are being parallelized to improve performance on large datasets. It is expected that this project will become the basis for a Master's thesis.

The parallelization of the projection step is done using P-threads under a shared memory processor architecture. Each sorting thread is added to a queue of consumers and one or more threads at the head of the queue are awakened each time the producer finishes a quantum of work. As each finishes, a final rendering thread converts the information stored by the projection threads directly into optimized OpenGL hardware calls.

Under a distributed memory architecture such as that of ASCI White and other supercomputing clusters, distribution of the data set in memory becomes very important. Therefore, the next step will be to load balance each dataset for optimal distributed behavior.

Curvature-based Adaptive Mesh Refinement for Level Set Tracking

Paul Covello
University of California, Davis

Summary:

Level set tracking is used to compute burn tables to be used for detonation shock dynamics codes. The level set represents the detonation shock front, which is assumed to propagate like an optical wavefront and thereby follows an eikonal equation. The advantage of using burn tables is that they greatly speed up and simplify the detonation shock dynamics code.

Algorithms to track level sets exist already. Given a grid, an initial level set, and a speed of detonation shock propagation, these numerical codes predict how the level set will propagate through the rest of the given domain. The advantages of contemporary methods are speed and versatility with unstructured grids. The disadvantage is that when the curvature of the level set is high, the accuracy of these algorithms falls off. There are two ways to remedy this: one is to have higher resolution, or more discrete points, where the curvature of the level set is high; the second is to provide more information to the grid node being updated about the local grid geometry.

I have developed two different algorithms pursuing these strategies. The first performs adaptive mesh refinement in regions of high curvature, and the second increases the discrete scope. Results show an increase in accuracy, sometimes attaining second order. The expected price was computational speed, sometimes taking up to two orders of magnitude longer in execution time. On balance, these algorithms would be appropriate for complex multiscale problems where obtaining similar high accuracy with conventional nonadaptive methods would require several orders of magnitude more execution time.

The results were from two-dimensional problems, both orthogonal and non-orthogonal, and structured, orthogonal three-dimensional problems. The next goal is to demonstrate both approaches on three-dimensional simulations involving both structured and unstructured grids.

Computational Modeling of a Solid Fuel Rocket Booster Using Overture

Nathan Crane

University of Illinois, Urbana-Champaign

Summary:

A Solid Rocket Booster (SRB) is an extremely complex system. Experimental data obtained to measure the performance of a prototype motor or determine the reason for failure of an SRB is both difficult and expensive to obtain. A good computational simulation of an SRB could help alert designers to potential problems without constructing the actual rocket, thus reducing the number of design iterations required. As the solid rocket fuel burns inside of the SRB, complex moving gas/solid interfaces are produced. This geometry is difficult to model with standard grid generation techniques. Overture is an overlapping grid software environment developed at LLNL and designed to combine the numerical and computational advantages of bulk structure with the geometric versatility of unstructured grids. The objective of my research is to investigate the ability of Overture overlapping grids to represent the complex moving geometry inside a burning rocket.

First, a 2D axisymmetric model of a SRB was produced. Incompressible steady state and transient compressible fluid flow equations were solved on the static 2D mesh. The effect of a moving geometry in 2D was also investigated.

Next a 3D computational model representing the full 3D geometry of the inside of a space shuttle SRB was produced. The model represented the flow space inside of a SRB with inflow of hot gas on the void boundary. As the rocket fuel burns the void space in the rocket will expand. The fuel-void interface velocity at a given point is a function of several variables. Propellant burn rate is a function of local pressure, flow velocity, and exposed surface area. In addition, structural deformation of the solid fuel can alter the shape of the void. The size of the 3D model produced (approximately one million grid points) makes time required for solution of the flow equations prohibitive. A grid regression scheme based only on exposed propellant burn area was implemented.

Overture-style overlapping grids proved to be very useful for modeling the complex geometry that may occur during the operation of a SRB.

Large Eddy Simulation of Rayleigh–Taylor Instability using the Arbitrary Lagrangian–Eulerian Method

Rebecca Mattson Darlington
University of California, Davis

Summary:

This research addresses the application of large eddy simulations (LES) to Arbitrary Lagrangian Eulerian (ALE) simulations of Rayleigh–Taylor instability. Large eddy simulation is a popular methodology for dispensing with the artifacts of turbulence modeling at the larger dynamical scales of the flow.

First, ALE simulations of a simplified Rayleigh–Taylor instability are produced. The advantages of ALE over purely Eulerian simulations are shown. Next, the behavior of the LES is examined in a higher-fidelity ALE simulation of the Rayleigh–Taylor instability. The effects of eddy viscosity and stochastic backscatter are examined. The LES is also coupled with ALE to increase grid resolution in areas where it is needed.

Finally, the methods studied above are applied to two sets of experimental simulations. In these simulations, ALE allows the mesh to follow expanding experimental targets, while LES represents the effect of unresolved instability modes.

Domain Decomposition for a Periodic Scattering Problem

Mike Flanagan
Texas A&M University

Summary:

I am pursuing domain-decomposed iterative techniques for wave-Helmholtz problems in two-dimensional periodic layered media. Solution data is passed between adjacent domains via a Dirichlet-Neumann operator. The goals are to establish convergence and determine means to improve convergence rate.

The domain is rectangular with periodicity in the x-direction. The layered structure exhibits parallelism since each structure is identical except for the incoming and outgoing wave data. The discretization on each subdomain is by finite elements with bilinear nodal basis elements. The algorithm employs a direct solver on each subdomain. In each iteration information passes to the next domain (above and below) by updating the outgoing wave and the incoming wave for each respective domain. This process was successful for four layers and this was the main test for working on shared memory systems (4 nodes), but as attempts were made to increase the number of layers, convergence was unreliable. Error decreased rapidly, but then started to diverge. Resonances are suspected.

A rigorous analysis of the mathematics is currently being undertaken. Simultaneously, we are exploring FFT-based discretization techniques to explicitly take advantage of the periodicity.

Data Mining with Sparse Grids

Jochen Garcke
University of Bonn, Germany

Summary:

Data mining is the process of finding hidden patterns, relations, and trends in large data sets. It plays an increasing role in commerce and science. Recently a new data mining technique for the classification problem was introduced by Garcke et al. The current work focused on expanding its capability and applicability for real data sets.

The technique is based on the regularization network approach but, in contrast to the other methods which employ ansatz functions associated with data points, we use a grid in a high-dimensional feature space for the minimization process.

To cope with the curse of dimensionality, we employ sparse grids. Thus, only $O(h \wedge(-1)n \wedge(d-1)))$ instead of $O(h \wedge(-d))$ grid points and unknowns are involved. Here d denotes the dimension of the feature space and $h = 2 \wedge(-n)$ gives the mesh size. To be precise, we use the sparse grid combination technique where the classification problem is discretized and solved on a certain sequence of conventional grids with uniform mesh sizes in each coordinate direction. The sparse grid solution is then obtained from the solutions on these different grids by linear combination. In contrast to other sparse grid techniques, the combination method is simpler to use and can be parallelized in a natural and straightforward way.

We looked into the possibility of applying this approach in the context of dimensional reduction, i.e., eliminating nonessential attributes which did not enhance classification. The results achieved on some test data sets look promising so far. Further we worked on changing the underlying discretization process, moving away from tensor-product based basis functions in d-dimensions to basis functions having d-dimensional simplexes as supports. This will allow the treatment of data sets with a few more attributes than with the old implementation.

Accelerating a Global Ocean Circulation Model

Aaron Herrnstein
University of California, Davis

Summary:

The climate and carbon cycle modeling group at LLNL uses models of ocean circulation to help in the study of human-induced changes on climate and of methods for potentially reducing those changes. As part of this research, we have begun working with a new model developed at LANL, the Parallel Ocean Program, or POP.

Existing ocean models, including the model developed at LLNL, use various numerical techniques to accelerate the convergence of the model solution to an equilibrium state. Two of these techniques are timestep-splitting and deep-ocean acceleration. Both of these techniques involve the rescaling/redefinition of time within particular model equations or geographic regions.

Without the use of these techniques the models take unreasonably long amounts of time to reach an equilibrium state, often more than two months of time on a massively parallel computer. Aggressive use of acceleration can reduce the time to equilibrium by more than a factor of 50. However, these techniques are largely untested in the POP model, so the objective of this work is to perform such testing.

To test the applicability of these acceleration techniques in POP, we have made a series of model simulations with different combinations of acceleration. The results of these runs will be compared to determine the specific impacts of these combinations on the equilibrium solutions.

The runs are accelerated toward equilibrium solutions by factors from 5 to nearly 100 faster than an unaccelerated model. More than 30,000 processor hours of an IBM SP computer were used—this corresponds to running a high-power desktop computer for more than 3 years and demonstrates the need for such acceleration techniques. Even so, more run time is still needed to bring one of the simulations to equilibrium.

Analysis of the data generated by these runs will be ongoing, concentrating on the potential distortions to the seasonal cycle and the impacts on the large-scale circulation, such as the poleward heat-transfer and strength of major currents. Also, more simulations will likely be needed to determine the best operational mode for the model.

Developing a Higher Order Godunov Method for Lagrangian Based Shock Hydrodynamics

Charles Hindman
University of Colorado

Summary:

Computational shock hydrodynamics is a difficult problem occuring in many fields, including aerodynamics, astrophysics, material deformations, and weapons design. We are investigating a new methodology for solving shock hydrodynamics problems by coupling a higher order Godunov method with a cell-centered Lagrangian scheme. The combination of the two allows the creation of a code that exactly conserves energy and momentum and does not require the use of an explicit artificial viscosity or hourglass filter, unlike typical staggered grid approaches.

Using the framework of A++ and Overture, software toolkits developed at Lawrence Livermore National Laboratory, a non-staggered grid, Lagrangian-based higher order Godunov solver was developed for the solution of the Euler equations in one and two dimensions. Various algorithmic options were included in this code, including velocity and energy updating methods and viscous damping, for testing and comparison purposes. The ability to set up different initial conditions and boundary conditions for various problems was included.

This code was tested on various common verification problems, including the Sod shock tube problem, the Woodward–Colella interacting blast wave problem, the 2-D Sedov problem, and the Noh problem in one and two dimensions. The results matched those reported in the literature for all cases except the 2-D Noh problem, which experienced significant mesh entanglement along the *x*- and *y*-axes. However, this is known to be typical of staggered grid Lagrangian methods as well. The next step will include comparison testing and speed-up analysis.

Adding Embedded Boundary Capabilities to SAMRAI

Jason Hunt
University of Michigan

Summary:

Engineering applications are often multi-scale. Instead of resolving the entire domain to the smallest scale, structured adaptive mesh refinement (SAMR) refines the mesh local to areas requiring higher resolution. LLNL's SAMRAI framework uses object-oriented techniques to provide a parallel SAMR framework for application developers. Engineering applications also involve complex geometry within the domain of interest and many interesting flow phenomena are driven by interactions with complex geometry. Therefore, it is very desirable to combine SAMR with complex geometry representation.

The current version of SAMRAI and a previous implementation of the embedded boundary technique, redistribution, were used as a starting point. The previous embedded boundary implementation was altered to be compatible with the current version of SAMRAI, and new data structures were introduced to enable more efficient algorithms to work with the embedded boundaries. We achieved embedded boundary capabilities for rectangular coordinates in two and three space dimensions. The code currently supports simple analytic geometry. With the framework that SAMRAI provides, we then combined the embedded boundary technique with AMR for rectangular coordinates in two space dimensions. Testing of the code was also performed to address conservation violations.

The re-redistribution technique, which ensures conservation near embedded boundaries that pass through differing levels of refinement, needs to be implemented, and work to enable parallel computations is required. In addition to extending the current work in rectangular coordinates to three space dimensions, there is interest to implement cylindrical coordinates in two space dimensions. Finally, the goal is to have a SAMR code that can handle general complex geometries.

The Validity of Paraxial Approximations in the Simulation of Laser-Plasma Interactions

E. McKay Hyde
California Institute of Technology

Summary:

High-intensity lasers such as those used in inertial confinement fusion produce high-density plasmas, which interact with the propagating light. Solving the Helmholtz equation to compute the laser scattering induced by variations in the plasma density remains too difficult for large computational domains. Hence, various paraxial approximations are often employed because of their efficiency. In this work we sought to establish the domain of validity of these paraxial approximations.

We compared the light intensities computed by solving the Helmholtz equation and five different paraxial equations on small computational domains and for a variety of density profiles. The Helmholtz equation was solved using a variant on the method recently described by Bruno and Sei and the paraxial equations were solved using the implementation described by Dorr and Garaizer of CASC. The density profiles were of two types: (a) small, analytically described combinations of peaks and valleys in a constant background density; and (b) small sections of density profiles computed in actual laser-plasma simulations.

Although it is difficult or impossible to describe a complete set of criteria by which one can determine if a given paraxial equation will be valid, we did observe a general trend. For background densities at and below 0.25 nc (nc is the critical density), the Helmholtz and paraxial equations all agree rather well. However, for background densities at and above 0.4 nc, we observe significant disagreement. Furthermore, no single paraxial approximation emerges as the best choice in all cases. Some features of the solution are captured well by one paraxial equation which are not captured well by another, and vice versa.

Analytical and computational comparison of the various paraxial equations continues. We hope to find ways to improve or at least predict solution accuracy without sacrificing too much efficiency.

Improved Parallel and Serial ILU Preconditioning

David Hysom
Old Dominion University

Summary:

Many scientific codes ultimately devolve to the solution of large, sparse systems of linear equations, $Ax = b$. The linear systems, whose solutions may consume a vast majority of the execution time, are frequently solved using preconditioned Krylov methods. Popular incomplete factorization (ILU) preconditioning, although popular and comparatively robust, is in general difficult to parallelize. Additionally, in many cases it is unclear how best to select the most effective runtime parameters (e.g., the fill level for ILU(k) or the dropping threshold for ILUT). We are investigating various shared-memory parallelization approaches and means of increasing preconditioner effectiveness while reducing preconditioner size.

A literature search reveals prevalent beliefs that preconditioners should be relatively small (have approximately the same number of nonzeros as the problem matrix, A), and that preconditioner formation (factorization) and application (triangular solves) should take little time, compared to the Krylov method. There is increasing evidence, however, that larger preconditioners (i.e., wherein more fill-in is allowed) are most effective at minimizing run time for certain classes of problems. One question investigated, therefore, was how to maintain the effectiveness that comes with increased (higher level) fill, while simultaneously reducing preconditioner size.

We experimented with several sparsification strategies: dropping small values from A before factorization; dropping small values from the factors; dropping with and without row-scaling, etc. For the systems studied, we found that dropping small values could reduce solution time by 10% or more, and dropping small values from A was somewhat more effective than dropping values from the factor. It remains unclear, however, how best to select dropping thresholds for production-code applications.

A second means of reducing preconditioner size is to compute and apply the preconditioner in single precision arithmetic. This strategy has the effect of reducing preconditioner size in terms of memory requirements, as opposed to nonzeros. Although slightly increasing iteration counts, single precision preconditioners were found to reduce run time by up to 20%.

I previously developed and published results for a parallel ILU (PILU) algorithm implemented using MPI, and showed scalability to several hundred processors. While at LLNL I became interested in a medium-scale, shared memory implementation. The motivation was to reduce solution time for a specific set of problems (e.g., 2D CFD filament problems of Petri Fast) produced by the LLNL's Overture code. These problems contain up to 400,000 unknowns, and were being solved sequentially, since Overture's current implementation is sequential. Using OpenMP, I implemented threaded versions of Block Jacobi and PILU preconditioners, and CG and BICGSTAB Krylov methods. Comparisons on the Compass Alpha cluster are indicate high promise for PILU. The resulting code is expected to be incorporated into a future release of Overture through the Hypre library.

Constrained ALE Grid Dynamics

Ana Iontcheva
University of California, Davis

Summary:

In transient structural dynamics simulations at LLNL, objects are typically modeled with unstructured hexahedral meshes. As the objects undergo compressible motion, the grid distorts according to predefined rules so that the accurate answers can be computed for the field equations while maintaining regularity in the grid connectivity. Often, the direction of portions of the grid must be forcefully constrained so that they do not venture into spatial regimes where they violate certain physical principles. For example, in contact-impact problems grid points from each object are constrained to remain in their physical space. The imposition of these constraints is performed using a Lagrange multiplier method and the grid points are moved via the solution of a matrix equation.

During the summer, the two-dimensional motion of a continuum over a solid object was studied. The continuum was modeled with a two-dimensional viscous Burgers' equation in Lagrangian coordinates. The grid motion was caclulated using the velocity of the medium. The original grid was defined using the True-Grid package. As predicted, grid distortion occurred very rapidly. Future work includes the use of a viscous grid motion equation with Lagrange multipliers to enforce the grid motion around the solid object.

Parallel Implementation of Mortar Finite Element Method in Three Dimensions with Multigrid Preconditioning

Chisup Kim
Texas A&M University

Summary:

The mortar finite element method is a nonoverlapping, nonconforming domain decomposition method in which each subdomain can be meshed independently. This provides a natural setting for parallel implementation. The resulting algebraic system, be it from parallel or serial implementation, requires efficient iterative methods and for this purpose, we construct a dual basis multiplier space and a multigrid preconditioner.

The continuity of the solution across the generally non-matching subdomain interfaces is imposed in a weak sense by using the multiplier spaces. At each such interface, a mass matrix has to be inverted. Our dual basis multiplier space consists of discontinuous functions and can be constructed in a simple fashion. Moreover, it produces diagonal mass matrices, which is essential to efficient computation of three-dimensional problems.

The mortar problem can be formulated as a nonconforming finite element method, which leads to a symmetric positive definite problem. The multilevel mortar finite element spaces are constructed based on successively refined meshes starting from a coarse one, and are nonnested. The multigrid algorithm considered is an extension of the two-dimensional case analyzed recently by Gopalakrishnan and Pasciak, which in turn is an application of the BPX method. We use the dual basis multipliers to further improve the efficiency. This produces a uniform preconditioner; that is the condition number of the preconditioned system is bounded independent of the problem size.

The parallelization was done using OpenMP, which is an interface for parallel programming using shared memory. This machine provides shared memory in the hardware. The memory is physically distributed across processors, but is accessible by all the processors.

In algebraic multigrid, one starts with a fine grid, matching or nonmatching, and obtains coarse grids by applying successively a coarsening strategy, e.g., element agglomeration. When coarsening independently in parallel in each subdomain, the resulting coarse grid will not in general match across subdomain interfaces. Furthermore, the elements at such interfaces will not necessarily have regular shapes such as triangles or rectangles. Mortar finite element method with dual basis multiplier spaces in this context will be studied and implemented to be run on a parallel machine.

Using Sapphire to Mine Astronomical Data

Imelda Kirby
University of Washington

Summary:

Data collection threatens to overwhelm our capacity to organize and review it for scientific information and inference. Project Sapphire in the Center for Applied Scientific Computing at LLNL is developing new techniques for analyzing massive quantities of data to extract features that are useful for scientists engaged in the research.

My project in data mining for astronomical data began with installing IRAF, and other software for astronomical images. I continued last summer's search for different automated ways of detecting asteroids in terabytes of astronomical data, this time focusing on SDSS (Sloan Digital Sky Survey) data. I studied the earlier successful application of decision trees to the classification of bent-double galaxies of the FIRST data set, while trying to develop some image processing techniques that would improve the results of the decision tree.

Algebraic Multigrid Homogenization

Stephan Knapek
University of Bonn, Germany

Summary:

Solutions for problems that model locally strongly varying phenomena on a micro-scale level require that all length scales present in the problem be resolved. However, due to storage requirements and numerical complexity, the grid for numerical simulation cannot be chosen fine enough to meet this requirement. Fortunately, in many practical applications the fine-scale details of the solution are not of interest, but only a coarse-scale solution is sought. Therefore, it is sufficient to work with averaged equations. The first step towards an accurate numerical treatment of these problems is the determination of an approximate mathematical model that captures the influence of the unresolved fine scales of the medium, and hence effectively disconnects the mesh size of the computational grid from the size of the heterogeneities. This "upscaling" or homogenization procedure results in equations with so-called effective coefficients that vary on a coarse scale. We investigate the use of algebraic multigrid coarse-grid operator constructions for upscaling and homogenization by interpreting Galerkin coarsening methods as discrete homogenization methods.

Our approach for the computation of a coarse-scale operator and an approximation to the effective diffusion coefficient is based on methods that have been used for some time in robust multigrid algorithms. Specifically, the coarse-grid operators that are constructed with matrix-dependent prolongations by means of the Galerkin approximation may be viewed as discrete homogenized operators. From the coarse-grid limit, operator approximations of the effective diffusion coefficient can be determined. However, we showed that a straightforward interpretation of the standard AMG coarse-grid operators does not necessarily lead to a reasonable homogenized operator. This is due to the choice of coarse grid points in the AMG method. Our current research is on constraints to place into the AMG-coarsening to lead to reasonable homogenized operators.

A Parallel Three-dimensional Magnetohydrodynamics Solver

Joseph Koning
University of California, Davis
Department of Applied Science

Summary:

The simulation of magnetohydrodynamics (MHD) requires the solution of fluid dynamics coupled with electrodynamics. An interesting problem in this field is a compact toroid moving through a plasma contained within a tokamak fusion device. An idealization of this problem is a superconducting sphere moving through a plasma. The sphere's motion will excite waves in the plasma. The energy loss through the interaction of the sphere with the plasma and magnetic field will result in the sphere slowing down and eventually stopping. This model system has been treated analytically by Newcomb. The superconducting sphere in an ideal plasma serves as a proof of concept for the problem of a compact toroid interacting with a tokamak plasma. This problem is three-dimensional and includes the fluid dynamical phenomenon of shocks, as well as electrodynamic phenomena such as Alfven and magnetoaucoustic waves.

I use an Eulerian fluid dynamics method coupled with a Vector Finite Element (VFE) method to simulate the sphere-plasma interaction. The Eulerian method treats the fluid description of the plasma and the VFE method treats the electric and magnetic fields. The vector basis functions of the VFE method can be used that maintain normal continuity for the magnetic field and tangential continuity for the electric field while maintaining zero divergence of the magnetic flux density.

Implementing the method involves combining a modem object-oriented fluid dynamics code with the object-oriented VFE code. Currently the parallel framework PETSc is used to integrate the VFE method that has been combined with the fluid dynamics code. A cross product operator has been implemented to accomplish this combination. The current framework has the ability to construct cross and dot products as well as the curl and div differential operators, for the vector basis, and grad operators for the scalar basis functions. Some preliminary fluid dynamics simulations of a sphere propagating through an uncharged fluid have been run. Future work will involve validation of the current framework and simulations of the sphere propagating through the charged fluid.

Development of Flexible Smoothers for Overture

Frank Koster
University of Bonn

Summary:

Smoothing operations are required at various places within LLNL's Overture code, for example in a multigrid scheme or as part of elliptic grid generation. Therefore, a class of flexible smoother functions should be designed for reuse in all the various tasks.

The objects a smoother acts upon are the coefficient matrix, the right-hand side, and the solution. These objects have to provide the smoother with functions for the defect calculation, the (optional) treatment of boundary conditions, and various special elements, e.g. the diagonal of the coefficient matrix. Through such an object-oriented approach, the smoother attains the desired flexibility.

In our design these functions are members of the supplementary class "CoefficientMatrix" which hides all details to the smoother. Presently, the class "CoefficientMatrix" can deal with second and fourth order operators in two and three dimensions, with general, cross, and constant stencils. Boundary treatment is presently provided for second order operators only.

For the smoother functions, there are two interfaces: "realMappedGridFunctions" and the low level "realArray." Furthermore, there are three ways of using the smoothers: "realMappedGridFunction" with inherited topology and boundary treatment; and "realArray" with or without inherited topology, and no boundary treatment. This provides the flexibility required for current and future adaptive mesh refinement extensions of Overture.

We will complete the present design with boundary conditions for fourth order operators and line smoothers.

Superresolution of Buried Objects in Layered Media by Near-Field Electromagnetic Imaging

Sean Lehman
University of California, Davis

Summary:

In the use of non-invasive waves of either electromagnetic or acoustic origin to probe layered media in near-field conditions, few researchers outside of optical microscopists have taken advantage of the evanescent part of the scattered field to enhance resolution. In this project, we have proposed and are developing an imaging technique to be used in near-field environments. Our technique achieves resolution beyond the diffraction limit, or equivalently "superresolution," by including the evanescent part of the field backscattered from objects buried in the medium.

Tomographic imaging is a collection of techniques to reconstruct images of the unknown internal structure of an object from fields transmitted through, and/or reflected from it. There are two widely used tomographic techniques: projection tomography and plane-to-plane backpropagation. The most widely used diffraction tomography technique is resolution-limited to approximately a half wavelength. We study the forward scattering process in order to develop a new diffraction tomography imaging technique which achieves resolution beyond the classical limit. We proved a new theorem which explains why the diffraction tomography method is resolution limited. We then derived a total field scattering relation, which includes both propagating and evanescent field components. We are developing a new reconstruction algorithm based upon the total field scattering relation. This full-field tomographic reconstruction algorithm includes both propagating and evanescent field components. We present reconstruction results from both simulated and real wide-band radar data which demonstrate that our new algorithm surpasses the resolution of most current techniques.

Alternating Tridiagonal Solvers for the Anisotropic Diffusion Equation Spatially Discretized by Linear Triangular Finite Elements

Linh Lieu
University of California, Davis

Summary:

Alternating Tridiagonal Solver methods here refer to generalizations of alternating-direction-implicit (ADI) methods taking advantage of direct tridiagonal system solution without regard to notions of spatial direction. Stable ATS methods for evolution/solution of the anisotropic diffusion equation on spatially nonuniform grids could significantly reduce execution times. Our goal was to test ATS methods on logically uniform, spatially nonuniform grids, using linear triangular finite-element (LTFE) spatial discretization.

ATS, Crank Nicolson (CN), and Backward Euler (BE) methods were implemented in our code as evolvers. Also, ATS was implemented as a solver (evolution to steady state). Whenever ATS was used, the spatial operator was additively split into tridiagonal matrices, each corresponding to a logical direction. For LTFEs on logically uniform triangular meshes, this splits the spatial operator into three tridiagonal matrices. These matrices were used in a Douglas-Gunn scheme to advance the diffusion equation by some small (evolver) or large (solver) time step. The Doss-Miller prescription was used for time step adaptation.

As an evolver, ATS was used with the FE lumped mass matrix. However the use of a consistent mass matrix with CN or BE improved accuracy. When stable, the ATS method is algorithmically efficient. However, stability has been achieved only on spatially uniform grids, and for LTFEs, only when all interior angles of the triangles are greater than 45 degrees. So far the method is conditionally stable on spatially nonuniform grids, so no general utility has been found. However, lumped mass ATS evolution on spatially uniform grids is sufficiently fast and accurate that it is worth using on appropriate large problems.

An Evolutionary Algorithm Library for General-purpose Optimization

David W. Littau
University of Minnesota

Summary:

Evolutionary algorithms are useful for solving optimization problems. They can be applied to any black box function with parameters which, when varied, produce outputs of varying quality. Unlike many traditional optimization techniques, no other information about the function is required.

Many techniques exist in the evolutionary algorithm hierarchy. It is desirable to have a general-purpose library of specialized functions that can be combined to both reproduce existing evolutionary algorithms and to experiment with new ones. Since evolutionary algorithms are inherently parallel, parallel techniques should be included.

The library was implemented in C++. Pure virtual base classes were used for the crossover, mutation, and selection classes to provide a common interface among the various techniques available to perform each operation. Inheritance from these base classes enforced the common interface, allowing the seamless integration of the different methods into a wrapper class. The wrapper class results in a framework in which a scientist who is interested in optimizing a particular function can do so without needing any knowledge of how evolutionary algorithms operate. A comprehensive collection of classes is available.

Experiments were performed with the code library, and some improvements in existing techniques were realized. The research with the library is ongoing. Furthermore, the library will be used to enhance the performance of some tools that are currently being used on other research projects.

Automated Scalability Performance Analysis of MPI Programs

Michael McCracken
Pennsylvania State University

Summary:

Scalability performance analysis of message passing parallel programs is a difficult task, especially as the size and complexity of programs and their respective trace files increases. Traditional forms of trace data analysis, such as visualization, can become difficult to use when analyzing very large parallel jobs, such as those being run by researchers at Livermore using the ASCI platforms. Scalability is an important concern for these large parallel codes; however, measuring it is difficult because one must analyze data from numerous performance experiments. Tools that give summary-style and specific prescriptive information are now necessary to understand the scaling behavior of large parallel programs.

We have developed MPSET, an innovative performance tool for MPI programs. This tool gives the programmer specific prescriptive information on the scaling behavior of components of their application, enabling analysis of large systems with significantly reduced effort. MPSET examines multiple trace files from runs at various numbers of tasks, and then shows MPI communication operations that scale poorly. MPSET calculates the mean duration of each MPI library call at each scale, and can then fit a curve for each call across the scales, highlighting curves that denote poor scalability or irregular performance. This allows the programmer to better focus their optimization effort. Results from the ASCI Compact benchmarks and the NAS parallel benchmarks demonstrate that MPSET can significantly simplify exploring the cause of scalability problems and, thus, improve the scalability of MPI applications.

Further development of MPSET is expected, with the possibility of tying it to a user-interface infrastructure in development at CASC, as well as adding useful analysis features to MPSET, such as automatic highlighting of MPI functions that scale poorly and possible work on automatic load imbalance detection.

Reengineering and Extending the DataFoundry Applet

Jason V. Morgan
University of Utah

Summary:

Prototype software is often developed without fully incorporating modern software engineering techniques. However, before the software is put into production use, the code is usually reengineered to utilize those techniques in an effort to reduce future maintenance costs. DataFoundry's Java applet, the graphical front-end to a collection of genomics data and tools, is starting to be used in genomics research and needs to be made easier to extend and maintain. New features also require introduction to improve the software's usefulness.

Originally, two large Java classes, DF_applet and GeneticsFrame, provided the bulk of the DataFoundry interface functionality. DF_applet defined the components required to form a query and served as the starting point for the interface. GeneticsFrame defined the interface for viewing and manipulating the results from a query.

New queries, with different input and coding requirements, were difficult to add to DF_applet. To address this problem, DF_applet was split into smaller classes. One of these manages the query interfaces, while the others provide the specific functionality required for each query.

Adding new features to GeneticsFrame was difficult and error-prone because of both its size and the complex interaction of internal components. This problem was addressed by identifying similar data structures and methods, then creating several generic classes to provide the same functionality. These new classes promote significant code reuse and thus are much easier to maintain. A second change, one that made extending the applet easier, was the addition of code supporting automatic menu creation through a distinguished class hierarchy.

The most prominent new feature added to the interface was the ability to view the results of a batch blast run. This required defining BatchBlastPanel, the query starting point, and BatchResultsFrame, which displays a high-level view of the results. The frame also provides several new data manipulation capabilities, including the ability to email data in a spreadsheet compatible format.

Reducing future maintenance costs required more than just good engineering techniques. During reengineering, the code was made more readable, and extensive documentation was written. Near the end of the summer, additional time was spent writing high-level documentation to guide future programmers.

While the redesign of the applet is complete, the DataFoundry team will continue to add new features in response to user requests.

Software Development for Sapphire

Dave Nault
University of Cincinnati

Summary:

We undertook several software engineering improvements on the Sapphire framework for automated scientific data mining.

We tested the Sapphire Basic Toolbox Classes, creating a new set of test programs and making changes to enhance usability and documentation.

We optimized the Histogram Classes by removed dependency on the Standard Template Library, creating a new base class to manage parallel processing and interval manipulation, and unifying the histogram class hierarchy.

We also created new classes to handle array indexing, resizing, and resource management.

We updated the Decision Tree code, changing the way nominal data is represented, now using numbers instead of strings. Five new splitting criteria classes were added, along with new code to use the new histogram classes. Together these changes lead to a 100% performance gain.

Finally, a new data viewer was created. This included defining a file format for images, modifying the FITS viewer to read files in the new format, and writing some glue code that allows any program to display a data object by writing it to a file and launching the viewer application.

Overlapping Grid Interpolation using the P++ Array Library

Stefan Nilsson
Chalmers University of Technology, Sweden

Summary:

A key ingredient when solving Partial Differential Equations (PDEs) using overlapping grid techniques is the interpolation of boundary values between component grids. As it is an essentially "unstructured" operation, its implementation on a parallel computer will differ substantially from that of the difference operators usually employed at the inner grid points.

We have implemented a small subset of the possible overlapping grid operators using the parallel array library P++. The serial version of P++ (A++) is presently used in the Overture project for overlapping grid computations. To implement the interpolation algorithm we had to temporarily leave the P++ framework and code the necessary message passing explicitly using MPI function calls.

Subgrid Scale Reaction Modeling for Application to Turbulent Reacting Flows

Diem-Phuong Nguyen
University of Utah

Summary:

My research involves developing a computational fluid dynamics (CFD) code that calculates both turbulent reacting flows and complex chemical kinetics. Practical reacting flow simulations are accomplished through the introduction of a subgrid scale (SGS) reaction model. My summer research at LLNL involved using Overture to incorporate a SGS reaction model to transport equations, bridging microscopic details to the macroscopic domain.

The Overture framework allows for simulation of physical processes in complex geometry. A general Overture code was written in C++ to solve a system of PDE's on a composite grid using multiple time-stepping techniques. Both Euler and Crank Nicholson time stepping were implemented. Several overture functionalities were exploited, including interactive plotting, time-step control based on the CFL number, and twilight zone forcing which uses the method of analytic solutions to generate an exact solution to the PDE.

Convection-diffusion-reaction equations for multiple species were tested. An SGS reaction model involving the Baldyga two-step reaction mechanism was used to provide the reaction source term in the species transport equations. The Baldyga reaction model calculates the chemical state space of the system given the independent mixing and reaction progress variables. The reaction rates are returned to the transport equations and the new chemical species concentrations are calculated. This coupling of the reaction model to the transport model is achieved via a Newton search algorithm.

The reaction computations were calculated on both single and composite grids. Comparison of the results yielded no difference between the two cases for the same geometry.

I plan to run more detailed flow simulations involving the Baldyga reaction model in Overture for validation purposes. I also plan to couple different and more complex reaction models to the code. Other models include equilibrium and a 6-step heptane mechanism, which includes soot chemistry.
Comparison of the different reaction models can then be made by running the different computations in the Overture framework using the same geometry. I will also test the system in a parallel environment.

2D and 3D Models of Rabbit Sinoatrial Cells Using Overture and CVODE

Chris Oehmen
University of Tennessee, Memphis

Summary:

Over 300,000 Americans die each year from sudden heart failure. Although some aspects of cardiac function have been studied for over one hundred years, the mechanisms by which cardiac fibrillation and defibrillation operate are still uncertain. Many investigators have succeeded in modeling various electrophysiological and metabolic aspects of cardiac cellular function, but more detail and computational power are needed to discern the interplay of cell- and tissue-level processes.

The present implementation of a rabbit sino-atrial node cell is based on the model and code of Demir et al., which was modified to contain additional current features not known until after the time of publication. For instance, the rapid and slow delayed rectifier potassium currents and the rapid and slow components of hyperpolarization current activation were both incorporated in the model. In addition, the model was developed to reflect the mean action potential (AP) characteristics observed by Verheijck et al (1998) for spindle cells.

The LLNL code Overture was used to simulate the cell in 2D and 3D environments. In the extracellular and intracellular spaces, simple diffusion was solved for K+ and Na+. Ca2+ was allowed to diffuse in the extracellular space, while it was uptaken and released from another interior space, the Sarcoplasmic Reticulum (SR), in accordance with the Demir model, as well as buffering in the SR and intracellular spaces.

The concentration gradient across the membrane determined the Nernst potential at all points on the membrane, for which the LLNL code CVODE was used to solve a system of nonlinear ODEs that determined the membrane currents. This membrane current activity was then allowed to interact with the diffusing species in the various spaces.

Overture allowed for rapid implementation and development of the models, as well as 2D, 3D, and movie-style visualization of the results. CVODE provided a robust mode of integration for variables on the membrane surfaces.

The models were in good agreement with those in the literature. Both the 2D and 3D cell models exhibited the critical aspects of rabbit SAN function, including endogenous pacing and the relative dominance of particular membrane currents during different phases of the action potential.

Future developments of the models include the incorporation of electric field calculations based on the charge distributions calculated by the models. This will be an integral part of the final models in understanding the behavior of the ions in the represented physiological systems.

If the models are further validated with respect to the processes of interest, they may be extended into a massively parallel implementation in which many cells can be modeled in 3D, and their tissue-level properties can be studied.

Scalable Domain Decomposition Algorithms for Resolving Contact Surfaces in ALE Computations

Tim Pierce
University of California, Davis

Summary:

In the Lagrangian approach to modeling material media, the overall volume of material under consideration is subdivided into a discrete set of zones, together with their connecting faces, edges, and nodes. The connectivity of the resulting mesh determines which nodes belong to each zone, from which can be deduced which zones contain a given node. This information establishes the proximity relations of all zones.

As the material moves in time under the influence of internal stresses, body forces and boundary conditions, the mesh follows the material, so that a given zone always contains the same material. For an explicit time advancement scheme, the behavior at a given node or zone over a single timestep is dictated only by material in its own and neighboring zones, which can be determined from the mesh connectivity. The timestep is chosen, with reference to the speed at which signals can propagate through the material (typically the sound speed), so that this locality of influence is a valid assumption.

If the model includes more than one discrete body, then the interaction of the bodies must also be considered, and mesh connectivity is no longer sufficient to determine proximity. The bodies may come together or separate, or may slide along each other, each exerting a boundary force on the other. Though the connectivity of the mesh representing each individual body is constant, which zones are adjacent across the contact surface is entirely dynamic and unpredictable.

A standard approach to parallelizing a Lagrangian dynamics code involves dividing the mesh into a number of submeshes, or subdomains, of approximately equal size, and assigning the subdomains to separate processors. A good decomposition minimizes communication, as only neighbor information is required. Furthermore, since the connectivity is constant in time the communication pattern can be set up during an initialization phase, and only the variable data itself updated each cycle.

This simple domain decomposition method of parallelization breaks down at contact surfaces, or at least an additional mechanism must be supplied to: (a) determine proximity relations across contact surfaces; (b) distribute the calculation of the contact forces across the parallel machine; and (c) gather the necessary data from across the machine together on the processor where that part of the surface is to be calculated. The task of calculating contact forces can be divided into two steps: contact detection and contact enforcement. First one must determine if and where two surfaces are in contact. Then one can apply suitable balancing forces at those points.

In a serial environment, the detection problem can always be solved by brute force. The position of each node on one side can be compared to the positions of all faces on the other side, the closest face determined, and then penetration checked for. In a parallel environment, however, even the brute force approach is not available, since the closest face on the other side may be on a different processor. As time changes, the closest face and perhaps the processor assigned to that face may also change. In fact, the parallelization of contact surfaces provides formidable challenges, particularly if the task must be done in a scalable fashion as the mesh size and processor count are increased proportionately.

In this research, we developed a truly scalable solution, implemented it in a major dynamics code (ALE3D), and demonstrated scalability well beyond 1,000 processors. A document describing the design and implementation of the algorithm in great detail is currently under development.

Shrink-wrapping Algorithms for Large Data Visualization

Serban D. Porumbescu
University of California, Davis

Summary:

Scientists are doing simulations that produce terabytes of data output. Ideally, scientists need to interact with and manipulate the data in real time. Unfortunately, contemporary computers have I/O systems that are out of balance with the rate at which they can create data. To achieve our goal of visualizing large amounts of data we need to invent new algorithms that take advantage of powerful wavelet compression schemes.

Given an unstructured surface, shrink-wrapping produces the topologically equivalent semi-structured mesh needed for wavelet compression. The algorithm begins with a coarse base mesh and repeatedly subdivides, smoothes, and snaps (moves towards the desired surface) until the mesh parameterization is complete. Our principal improvement to shrink-wrapping involved modifying the smoothing step. The new smoothing operation performs edge-length-weighted Laplacian smoothing. This new technique greatly improves the output mesh by helping to prevent mesh vertices from gathering towards one particular part of the surface being parameterized. This work along with other related results was presented at the NSF/DoE Lake Tahoe "Workshop on Hierarchical Approximation and Geometrical Methods for Scientific Visualization" in October 2000.

Certain features remain difficult to obtain during shrink-wrapping because of the coarseness of the initial base mesh. We are currently investigating new techniques to ensure that shrink-wrapping can capture all of the features of the original unstructured mesh.

Finite Element Solutions to Laplace's Equation

Robert N. Rieben
University of California, Davis

Summary:

The purpose of this work is to develop a finite element program using the programming language Mathematica that could solve Laplace's equation with second-order accuracy using linear, tetrahedral finite elements. This program will serve as an implementation base for more general problems.

The program was developed from the ground up using the symbolic algorithms of Mathematica. The basis functions for the approximate solution are constructed using tetrahedral shape functions defined explicitly for each element (as opposed to being defined with respect to a reference element). The test problem is a unit sphere with Dirichlet boundary conditions. To demonstrate convergence, the solution was approximated on a series of ten successively refined meshes and compared to the exact solution. Plotting routines were developed.

Future additions to the program include the use of other three-dimensional elements (hexahedrons, prisms, etc.) and higher order finite elements. A time stepping routine will be developed for time dependent equations. This will lead to an adaptive mesh routine, which will restructure the elements of the mesh on the each time step based on the evolving solution.

A Multigrid Strategy for Accelerating Steady-State Computations of Waves Propagating with Curvature Dependent Speeds

Jonathan Rochez
University of California, Davis

Summary:

During the past year, we have been developing a code for steady state solutions to the eikonal equation in 2D and 3D using two differencing schemes for the multigrid method. Standard iterative methods show a quick reduction of the residual followed by a slow final convergence to the solution at high iterations. Such systems are ripe for the use of multigrid methods to speed up convergence.

Numerically, two different approaches have been investigated. The first approach organizes the calculation point by point on the grid using a finite-differencing scheme. The advantages of this approach include its speed in calculation time and ease of implementation, while its disadvantage is restriction to regular orthogonal grids. During this past year, code was developed to solve the eikonal equation on any orthogonal geometry with one or several detonation points. Curvature-based speeds for the wavefront propagation have been successfully implemented.

On a simple square grid geometry, solution rates of several sizes of problems with several multigrid V-cycles have been compared. In all cases, the speed was significantly faster for the multigrid method compared with calculation on a single grid. The finer the resolution, the greater the speed-up observed. For the simple geometry, the analytical solution of the problem is known to be circles emanating from each detonation point. From analyzing Fourier spectra of the error, we know it to be smooth throughout the calculation, implying the well suitedness of multigrid.

The second approach is based on discontinuous Galerkin method finite elements in a zone by zone calculation. The advantages include unstructured grids for the formation of non-regular geometries and fewer data points, while the disadvantages include larger calculation times and a less straightforward implementation. No results have been obtained yet since this second code is still being developed.

New Means of Compressing Images Using Wavelets

Joshua Senecal
University of California, Davis

Summary:

As computers continue to become more powerful the simulations run on them become larger. The result is the amount of data generated by these simulations—already on the order of terabytes—is also increasing. With the specific goal of storing the PPM dataset from the Richtmyer–Meshkov execution that won the 1999 Gordon Bell Best Performance Prize, we are developing a lossless compression method that achieves better compression than anything currently available.

Our method for compression involves transforming the data by a linear wavelet transform, and then compressing the resulting coefficients. Two observations about wavelet coefficients are these: most of the values lie at or near zero, and all the bits after the coefficient's leading bit position are random in appearance. We exploit this by compressing the position of the leading bit of the coefficient's magnitude. All zeros before this position are ignored, and all data after the leading position are copied as is. To encode the leading bit position we use Huffman encoding. Huffman's original algorithm generates codes based on the global probability of any particular symbol occurring. In our program we use context to improve compression. For example, if our program is currently processing a coefficient with a leading bit position of 1, the next coefficient is more likely to have a leading position of 0 or 1, and less likely to have a leading position of 7 or 8. We therefore construct Huffman tables based on these conditional probabilities.

As an example of how we encode, say that we need to compress the coefficient –5. In binary it is represented by a sign bit followed by an 8-bit magnitude: 100000101. Assuming that we already have a Huffman table generated, encoding would proceed as follows:

- Identify the leading bit position of the magnitude, which is 3 in this case.
- Look up the corresponding binary Huffman code for that position, say 10 in this case.
- Write out the encoded coefficient by first writing the position code (10), the sign bit (1), and everything that followed the leading bit position (01). The encoded coefficient is therefore 10 + 1 + 01 = 10101. The coefficient has been reduced to 5 bits.

As our benchmark we use the common compression program gzip. Four images, converted to 8-bit grayscale, were used in testing our compression program. One was a photograph, one was a black-and-white comic, and the other two were scientific visualization images. When compressing the photograph our program beat gzip, but gzip gets better compression on all the other images. Other tests show that our program appears to consistently do better than gzip when compressing images (like photographs) that have a lot of variation in them. gzip does better if the images have a small number of colors, contain a lot of regularity, or have large areas containing one color.

This program is still in its early stages, and more improvements are needed. Future work on improving the compressor will involve investigating additional ways of generating better Huffman codes based on context, and on improving our method so that it works better on a wider range of input

Parallel Algebraic Multigrid Methods

Marc Alexander Schweitzer
University of Bonn

Summary:

Multigrid methods are known to be fast solvers for elliptic PDEs. Specifically for strongly elliptic symmetric PDEs of second order, optimality of the convergence rate is guaranteed under weak regularity assumptions or even without regularity assumptions for the multiplicative and additive forms of the algorithm. However, for more general elliptic problems with anisotropic or discontinuous coefficient functions, standard multigrid methods using linear interpolation, full weight restriction and full coarsening lead to poor convergence rates, which are moreover dependent on the coefficient functions. Therefore, an asymptotic estimate for the condition number of the preconditioned system may be only of limited use in applications. This missing robustness has led to modifications of the standard multigrid method, based on modifications of the interpolation and restriction operators, use of multiple coarse grids, specially adapted coarse grids, or combinations of these techniques. When the interpolation operators are chosen dependent on the fine grid matrix (in a specific way) then in one dimension multigrid is just cyclic reduction, i.e., a direct solver, and in two dimensions (together with an ILU-type smoother) a robust method results. The so-called algebraic multigrid method (AMG) due to Ruge and Stüben also belongs to this class of multigrid methods, but allows for additional flexibility, through the choice of coarsened grids in an operator dependent way. Then a simple pointwise smoother is sufficient for a robust convergence behavior for a large class of problems. The basic idea of an algebraic multigrid (AMG) solver is the use of operator/matrix dependent coarsening and interpolation to achieve robustness of the solver.

The coarse grid selection within the AMG by Ruge and Stüben is inherently sequential. Numerous parallel coarse grid selection algorithms have been proposed to overcome this problem. Most approaches try to fix this problem by localizing the coarse grid selection to local process data and applying an additional setup phase to match the resulting coarse grids at process boundaries. The quality of the coarse grids generated by these approaches are dependent on the quality (with respect to AMG) of the underlying data partition of the linear system. Not only is the quality of the coarse grids affected by the data partition, but the scaling behavior of these approaches may be significantly limited by data partitions, which do not respect the physics of the underlying PDE and the discretization process.

Summary (continued):

We have proposed applying parallel multilevel graph-partitioning heuristics to the partitioned linear system to recapture the physics in the system and re-partition the data accordingly. The reduction in run time by applying AMG on an appropriate data partition though may be too small to justify this re-partitioning step by default since this pre-setup step also comes at a certain cost. For instance data partitions coming from classical domain decomposition approaches and discretizations on uniform grids and constant coefficient problems already capture the physics of the underlying problem and therefore a re-partitioning of the system will not improve the quality of the partition significantly. Modern adaptive discretization methods (adaptive FEM, meshless methods, etc.), however, usually provide the data partition either by using graph-partitioning techniques or even by some other spatial ordering technique (space-filling curves, etc.) which lead to more general data partitions.

Hence, the design of a cheap automatic mechanism to determine the 'quality' of a data-partition is necessary. The first step though towards such an algorithm is an algebraic definition of 'quality' for this purpose. To do so we have to study the effect of 'ill-partitioned' data on the quality and scalability of AMG. Here, the test problems for AMG have to be complex enough. We are working on a parallel port of a generalized (three-dimensional non-uniform discretizations) test suite to study the quality of (sequential, additive) AMG on general anisotropic diffusion and convection-diffusion problems. Furthermore, we are working on a parallel meshless method using space-filling curves for the data-partition and the interface to the hypre library.

Modeling Protein Production and Inhibition from DNA Regulation

Bahrad Sokhansanj
University of California, Davis

Summary:

This year marked the successful completion of the Human Genome Project and other whole-genome sequencing projects for rice, the fruit fly, and countless microbes. The result of these projects is the DNA sequence of almost all the genes of these organisms, and the sequences of all the protein molecules that their cells can produce. There are also small sequences of non-coding DNA which provide sites for protein binding and promote or inhibit the gene activity. However, this information is akin to a parts list for an electrical circuit: there is no wiring data, and because of limits to our current biological knowledge the function and quantity of the parts is also unknown. A major effort in bioinformatics has been to computationally predict the function of all the proteins and the wiring diagram of protein and DNA interactions. But, living cells are dynamic, constantly changing the kinds and quantities of protein molecules that they produce in response to environmental conditions: static information needs to be applied in kinetic simulation.

Our goal is to develop a dynamic simulation of bacterial gene regulation. Regulation in humans and other higher organisms uses the same fundamental processes, while adding more complex details. The model will include the binding of inhibition and activation proteins to DNA, the transcription of message RNA from the gene, and translation of the RNA to its protein product. Recent work has shown that this process depends heavily on stochastic effects resulting from the small number of molecules involved, so the traditional deterministic chemical kinetics approach is inadequate. Molecular dynamics is unsuitable for long time scales, so we are applying Monte Carlo methods with discrete numbers of molecules. Data from the literature and experiments at the Biology & Biotechnology Research Program (BBRP) are used for parameters and model validation. The long-range goal of this project is to dynamically model the genetic regulation of virulence in pathogens, and hopefully lead to a model of a whole bacterial cell.

ALE Computations Using Augmented Lagrangian Grid Speeds

Jay F. Thomas
University of California, Davis

Summary:

The Arbitrary Lagrangian–Eulerian (ALE) concept has been successfully applied to the simulation of a variety of continuum flow problems, including impact phenomena, material processing, and fluid-structure interaction. The ALE concept combines the Lagrangian approach, in which the computational mesh moves with the material, with the Eulerian approach, wherein the mesh is fixed and the material flows through it. The Lagrangian approach simplifies application of boundary conditions and tracking of material interfaces, but some flows may severely compress and distort the mesh. This causes the time step used to advance the solution to become extremely small and reduces solution accuracy in the distorted region. The Eulerian approach allows arbitrarily large deformations and turbulent flow, but material interfaces and certain types of boundary conditions may require special accommodation. The ALE concept attempts to circumvent these difficulties and obtain the advantages of both approaches by allowing the mesh to be moved in an arbitrary manner. The issue becomes one of specifying a mesh motion satisfactory for accuracy that can be efficiently generated.

Methods for specifying mesh motion can roughly be classified as static or dynamic. Static methods rely on computing the grid point locations at any time level by solving a steady grid equation (e.g., elliptic grid generation), and then interpolating the solution from the old mesh to the new mesh at a fixed time. Dynamic methods, by contrast, directly evaluate grid velocities via any appropriate law. The grid velocity equation and the physical equation may be solved simultaneously for the physical solution and the mesh. Interpolation of dependent variables from the old mesh to the new mesh is unnecessary.

ALE implementations have typically used a static approach to mesh motion as typified by mesh relaxation. Mesh relaxation is based upon elliptic mesh generation methods. The coordinates of the nodes are formulated as the solution to an elliptic problem and Jacobi iteration is used instead of a direct solution method. This results in a stencil that can be applied efficiently over the mesh. Only a few sweeps of Jacobi are used at each time step. The method thus indirectly specify the mesh node velocities but it does not lend itself to analysis in this regard.

The focus of the current effort is to investigate methods for directly specifying grid velocities in an ALE context. The general form chosen for the grid velocity consists of the fluid velocity augmented by an additive “correction” term. The idea is that the fluid velocity provides general solution adaptivity while the correction term can be chosen to control, for example, excessive mesh compression. Computations for Burgers’ equation and the 1D Euler equations, using Godunov’s method for the physical equation, have been performed using a correction term which measures how well a given weight function is equidistributed over the mesh. The approach appears effective for addressing mesh compression. To date a constant weight function has been used. Obvious extensions include use of a velocity-dependent weight function. The resulting grid velocity equation is parabolic. A significant issue for extension of the method to multiple dimensions is efficient solution of this equation, since it must be solved at each time step and stability considerations indicate that an implicit solution technique be used. The next phase of the investigation will address the extension to multiple dimensions.

ROSE Meta-program for Generating Optimizing Preprocessors

Danny Thorne
University of Kentucky

Summary:

The Overture object oriented (C++) framework for solving PDEs is hard for a compiler to optimize because it is designed with many high-level abstractions. Compilers are not capable of understanding the semantics of these high-level abstractions. ROSE is designed to be taught about the semantics of these high-level abstractions by way of specialized grammars corresponding to the code to be optimized. Thereby, ROSE will be able to generate a preprocessor that is capable of introducing optimizing transformations involving the high-level abstractions. The resulting source code, after the transformations applied by the ROSE generated preprocessor, should then be such that a compiler can build more efficient target code. One of the primary examples of the application of ROSE in Overture is to the array class library, A++/P++, from which all of Overture is derived. This library makes the develpement and use of Overture much more convenient than would be possible with primitive language constructs. However, the high-level abstractions that give rise to this convenience for the developer also have the effect of disabling the compiler from producing efficient target code. ROSE will be able to automatically transform many, if not all, uses of the array classes in a given code into something semantically equivalent but based on primitive language constructs and, hence, compiler-friendly. Since essentially all of Overture is built from the A++/P++ library, applying ROSE to transform the uses of A++/P++ classes in this way will provide for much more efficient Overture codes.

One shortcoming of ROSE, currently, is that it doesn't handle C++ templates. Many C++ codes today use templates. Hence, it is important for ROSE to be able to support templates. My main objective has been to devise and explore potential solutions to the problem of supporting templates in ROSE and to implement the support of templates if possible.

Three approaches for handling templates in ROSE have been attempted. The first two had to be discarded after much experimentation and analysis. The third is under development and promises to be successful. ROSE will be developed into a general-purpose tool for automatically introducing optimizing transformations involving the high-level abstractions in an object oriented (C++) code.

Error Control and Parallel Adaptive Grid Refinement for Convection-Diffusion-Reaction Problems

Stanamire Z. Tomov
Texas A&M University

Summary:

Simulation of flow and transport in porous media gives rise to extremely large-scale computations, making the consideration of iterative solvers natural. The solutions of such problems exhibit local nonsmooth behavior. Therefore it is essential that local grid refinement, based on a posteriori error analysis, be applied. Also, fundamental physical limitations on the computer processing speed may require the exploitation of parallelism. The goal is to create a tool that is based on discretization techniques utilizing finite elements/volumes, efficient preconditioning (in parallel) of the resulting sparse system, error control, and adaptive grid refinement.

Our computational approach is described as follows. We use mesh generator (Triangle for 2D meshes and NETGEN for 3D meshes) to generate a good coarse mesh. Then the problem is solved redundantly on the coarse mesh by every processor. The solution is used to compute a posteriori error estimates, which are used as weights in an element-based splitting of the coarse mesh into subdomains, using MeTiS. This splitting insures that the local refinements that follow will produce a computational mesh with number of triangles or tetrahedra balanced over the subdomains. Every subdomain is mapped to a processor. Then, based on a posteriori error analysis, each processor refines its region independently. After every step of independent refinement there is communication between the processors in order to make the mesh on that level globally conforming.

Concerning the a posteriori error analysis, I worked with Dr. Raytcho Lazarov on article "Error Control and Adaptive Grid Refinement for Convection-Diffusion-Reaction Problems in 3D." The article contains the description of an adaptive numerical technique based on finite volume approximations and the computational results of various model simulations of steady-state single phase flow and transport of passive chemicals in nonhomogeneous porous media in 3D.

I developed a 2D code with functionality as described in the computational approach above. The multilevel structure is used to define multigrid preconditioners. I worked closely with Panayot Vassilevski and Charles Tong on connecting this software to the Hypre Preconditioner Library. The connection is implemented using The Finite Element Interface (FEI) specification, which provides a layered abstraction that minimizes concern with internal details in the Hypre library. The initialization is done in parallel. GUI, using Motif, has been developed to utilize the selection of different Hypre options. After the solution is obtained on certain level it is sent directly through the AFJNET socket to the visualization tool SG, which resides on the user's machine. The benefit in this strategy is that the parallel machine (usually remote) is used only for computations and the local machine handles the visualization. This idea is very efficient for real time visualization of time dependent problems. The 3-D version of the code has the same features. Still under construction is the parallel local refinement step of communication between the processors for making the mesh on certain level globally conforming.

Raviart-Thomas zero-order (RTO) finite elements have been added to both the 2D and 3D codes. Under development is discontinuous approximation of convection terms for mixed finite elements. Both the 2D and 3D codes are written in C++. The parallel computations are done using the Message Passing Interface library (MPI). I finished with Raytcho Lazarov and Panayot Vassilevski article on "Interior penalty discontinuous approximations of elliptic problems," which has been submitted to *SIAM J. Scientific Computing*.

Interactive Visualization of Three-Dimensional Adaptive Mesh Refinement Data

Kevin Vlack
University of Illinois

Summary:

Adaptive Mesh Refinement (AMR) is a popular approach in scientific computing for focusing computational resources on regions of a problem where small step sizes in time and space are necessary for accurate calculation. While AMR is a natural solution for efficiently integrating stiff PDEs, it produces complicated datasets which are difficult to analyze without the aid of customized visualization packages. The main objective of the summer was to design and implement an efficient and practical volume-rendering module to interactively visualize 3D datasets created with an AMR algorithm.

There are many inherent characteristics of AMR that complicate the volume rendering process. A problem domain under an AMR framework is represented as a hierarchy of overlapping grids of increasing resolution, and each level of refinement completely contains the next, that is, any particular region of a refined grid is also represented within a coarser grid. Visualization entails selecting a mutually exclusive subset of grids from the data, and presenting the data in an interpretable manner. The situation also arises in which the scientist wishes to view overlapping areas in numerous refinement levels in order to analyze the relationship between them.

The volume renderer was implemented for X Windows in C++, using the Motif and BoxLib APIs. BoxLib is an AMR framework developed by the CFD group at Lawrence Berkeley Laboratory and is currently in use by researchers in both LBL and LLNL. The renderer uses the "splatter" technique, which approximates a voxel as an orthogonally projected polygon in order to speed the rendering process. It has been integrated into a prototype version of PAMRVIS, a parallel AMR visualization tool developed by Vince Beckner at LBL and expanded upon by Vlack in 1999.

The volume renderer was able to perform at interactive frame rates for numerous available AMR datasets, including a representation of a shear layer mixing and of a jet stream vortex. Developments were also made to the GUI of the application in order to intuitively change the parameters of the visualization in order to visualize various qualitative traits of the dataset. Although time constraints limited in-depth experimentation, basic timing techniques revealed that the maximum speedup achieved from parallel processing was about 2.8 when using four processors.

Future plans include integrating the volume-rendering module and additional interface features into the working version of the parallel visualization tool, PAMRVIS, which is used throughout the BoxLib community.

A Reduced Grid Method for a Parallel Global Ocean General Circulation Model

Michael Everett Wickett
University of California, Davis

Summary:

A limitation of many explicit finite-difference global climate models is the timestep restriction caused by the decrease in cell size associated with the convergence of meridians near the poles. A computational grid in which the number of cells in the longitudinal direction is reduced toward high-latitudes, keeping the longitudinal width of the resulting cells as uniform as possible and increasing the allowable timestep, is applied to a three-dimensional primitive equation ocean-climate model. This "reduced" grid consists of subgrids that interact at interfaces along their northern and southern boundaries, where the resolution changes by a factor of three. Algorithms are developed to extend the finite difference techniques to this interface, focusing on the conservation required to perform long time integrations, while preserving the staggered spatial arrangement of variables and the numerics used on subgrids. The reduced grid eliminates the common alternative of filtering high-frequency modes from the solution at high-latitudes to allow a larger timestep and reduce execution time per model step by roughly 20 percent.

Our reduced grid model is implemented for parallel computer architectures with two-dimensional domain decomposition and message passing, with speedup results comparable to those of the original model. Both idealized and realistic model runs are presented to show the effect of the interface numerics on the model solution. First, a rectangular, mid-latitude, flat-bottomed basin with vertical walls at the boundaries is driven only by surface wind stress to compare three resolutions of the standard grid to reduced grid cases, which use various interface conditions. Next, a similar basin with wind stress, heat, and fresh water forcing is used to compare the results of a reduced grid with those of a standard grid result while exercising the full set of model equations. Finally, global model runs, with topography, forcing, and physical parameters similar to those used for ocean-climate studies, are advanced to a near equilibrium state for both the reduced grid and the standard grid. Differences between the two are presented for typical fields of interest, and very little degradation of the solution due to the reduced grid is observed.

Integrating Arches with PETSc

Wing K. Yee
University of Utah

Summary:

Arches is a Computational Fluid Dynamics code being developed by Dr. Phil Smith and his colleagues at the University of Utah. The motivation for the development of this code is to simulate large-scale fire. First-principles simulations of fire are extremely complicated, taking into account mixing, reaction, and radiation, in addition to complex fluid mechanics. For large-scale fire, not only are the variables resolved on the grid scale important, the subgrid level variables are also important and need to be modeled. The resulting multiple-scale nonlinear partial differential equations need to be solved accurately on each time step. DOE's PETSc (the Portable, Extensible Toolkit for Scientific computation) is a software environment with solvers for discretized PDEs in parallel or serial. It has been employed to port legacy CFD codes to the ASCI parallel platforms, and one such port resulted in a 1999 Gordon Bell Prize. Flexibility is a trademark of PETSc. It has a Newton method base solver, but offers many different options for preconditioning and matrix construction. It is also compatible with all three major high-level languages: FORTRAN, C, and C++. The price of working with PETSc is allowing it to control the major distributed data structures, and supplying routines that link the physics-laden user code in local (legacy) data structures with the PETSc data structures.

Our first step to integrate PETSc into Arches involved an older version of Arches, Banff. This has provided valuable guidance for the development of Arches. From this learning experience, a programming standard for the Arches code has been set. The tests with Banff were done on the ASCI Blue machine.

In addition to the parallelization project, a new Monte Carlo numerical integration method for the mixing model was also explored.

Scalable Adaptive Multilevel Algorithms

Gerhard Zumbusch
University of Bonn

Summary:

The solution of boundary value problems of partial differential equations is considered. The goal is to construct an efficient solver, which uses an optimal order solution algorithm (multigrid), a low number of unknowns (adaptivity), very little memory (hash storage), and runs on large parallel computers. It is mainly the grid adaptivity that poses some difficulties for the parallel implementation.

Multigrid and multilevel methods are optimal order solution algorithms for equation systems stemming from the discretization of PDEs. They require linear time, that is O(n) operations for n unknowns. In the presence of singularities of the solution the convergence of standard discretizations is not as rapid as for smooth, regular solutions. One way to overcome this problem is adaptivity. The grid is refined during the computation, where indicated by an error estimator or error indicator.

The components of our adaptive multigrid code have to be parallelized. A single data decomposition is used on the distributed memory computer. However, each component requires a different treatment. A dynamic load-balancing scheme is used for the parallel execution of the code. It is based on space-filling curves. The parallelization technique has also been applied to sparse grid discretizations.